

**APPLICATION FOR
UNITED STATES PATENT**

in the name of

Shuwu Wu

of

America Online, Inc.

for

Voice Instant Messaging

Fish & Richardson P.C.
601 Thirteenth Street, NW
Washington, DC 20005
Tel.: (202) 783-5070
Fax: (202) 783-2331

ATTORNEY DOCKET:

06975-099001

Voice Instant Messaging

This application claims the benefit of U.S. Provisional Application No. 60/189,974 filed March 17, 2000 and U.S. Provisional Application No. 60/239,917 filed October 13, 2000.

TECHNICAL FIELD

5 The present invention relates generally to transferring data between subscribers of a communications system and more particularly to transferring audio data between subscribers of an instant messaging host.

BACKGROUND

10 Online service providers are constantly offering new services and upgrading existing services to enhance their subscribers' online experience. Subscribers have on-demand access to news, weather, financial, sports, and entertainment services as well as the ability to transmit electronic messages and to participate in online discussion groups. For example, subscribers of online service providers such as America Online or CompuServe may view and retrieve information on a wide variety of topics from servers located throughout the
15 world. A server may be maintained by the service provider or by a third party provider who makes information and services available through the worldwide network of computers that make up the online service.

20 America Online has provided subscribers with the ability to send and receive instant messages. Instant messages are private online conversations between two or more people who have subscribed to the instant messaging service and have installed the necessary software. Because such online conversations take place in essentially real time, instant messaging can provide immediate access to desired information. Instant messaging is becoming a preferred means of communicating among online subscribers.

SUMMARY

In one general aspect, electronic data is transferred between users of a communications system by enabling instant messaging communication between a sender and at least one recipient through an instant messaging host. In addition, voice communication is enabled between the sender and the recipient through the instant messaging host.

Implementations may include one or more of the following features. For example, implementations may include receiving and authenticating a text instant message from the sender at the instant messaging host; determining capabilities of the recipient; reporting the capabilities of the recipient; receiving a request to establish voice communication from the sender and/or the recipient; and/or authenticating the request. Authenticating may include identifying a screen name and/or an IP address of the sender and/or the recipient. Determining capabilities of the recipient may include identifying hardware or software associated with the recipient. A user interface may be displayed according to the capabilities of the recipient.

Voice communication may be enabled by establishing a generic signaling interface channel, a control channel, and an audio channel between the sender and the recipient. A mode UDP test may be attempted on the audio channel. The control channel may include a TCP/IP socket. The audio channel may include a UDP or TCP channel.

These and other general aspects may be implemented by an apparatus and/or by a computer program stored on a computer readable medium. The computer readable medium may comprise a disc, a client device, a host device, and/or a propagated signal.

Other features and advantages will be apparent from the following description, including the drawings, and from the claims.

DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a communications system.

Figs. 2-5 are expansions of the block diagram of Fig. 1.

Fig. 6 is a flow chart of a communications method that may be implemented by the systems of Figs. 1-5.

Figs. 7-10 are illustrations of different graphical user interfaces that may be provided by the systems of Figs. 1-5.

DETAILED DESCRIPTION

For illustrative purposes, Figs. 1-5 describe a communications system for implementing techniques for transferring electronic data. For brevity, several elements in the figures described below are represented as monolithic entities. However, as would be understood by one skilled in the art, these elements each may include numerous interconnected computers and components designed to perform a set of specified operations and/or dedicated to a particular geographical region.

Referring to Fig. 1, a communications system 100 is capable of delivering and exchanging data between a client system 105 and a host system 110 through a communications link 115. The client system 105 typically includes one or more client devices 120 and/or client controllers 125. For example, the client system 105 may include one or more general-purpose computers (e.g., personal computers), one or more special-purpose computers (e.g., devices specifically programmed to communicate with each other and/or the host system 110), or a combination of one or more general-purpose computers and one or more special-purpose computers. The client system 105 may be arranged to operate within or in concert with one or more other systems, such as for example, one or more LANs ("Local Area Networks") and/or one or more WANs ("Wide Area Networks").

The client device 120 is generally capable of executing instructions under the command of a client controller 125. The client device 120 is connected to the client controller 125 by a wired or wireless data pathway 130 capable of delivering data.

The client device 120 and client controller 125 each typically includes one or more hardware components and/or software components. An example of a client device 120 is a general-purpose computer (e.g., a personal computer) capable of responding to and executing instructions in a defined manner. Other examples include a special-purpose computer, a workstation, a server, a device, a component, other equipment or some combination thereof capable of responding to and executing instructions. An example of client controller 125 is a software application loaded on the client device 120 for commanding and directing communications enabled by the client device 120. Other examples include a program, a piece of code, an instruction, a device, a computer, a computer system, or a combination thereof, for independently or collectively instructing the client device 120 to interact and operate as described herein. The client controller 125 may be embodied permanently or

temporarily in any type of machine, component, equipment, storage medium, or propagated signal capable of providing instructions to the client device 120.

The communications link 115 typically includes a delivery network 160 making a direct or indirect communication between the client system 105 and the host system 110, irrespective of physical separation. Examples of a delivery network 160 include the Internet, the World Wide Web, WANs, LANs, analog or digital wired and wireless telephone networks (e.g. PSTN, ISDN, or xDSL), radio, television, cable, satellite, and/ or any other delivery mechanism for carrying data. The communications link 115 may include communication pathways 150, 155 that enable communications through the one or more delivery networks 160 described above. Each of the communication pathways 150, 155 may include, for example, a wired, wireless, cable or satellite communication pathway.

The host system 110 includes a host device 135 capable of executing instructions under the command and direction of a host controller 140. The host device 135 is connected to the host controller 140 by a wired or wireless data pathway 145 capable of carrying and delivering data.

The host system 110 typically includes one or more host devices 135 and/or host controllers 140. For example, the host system 110 may include one or more general-purpose computers (e.g., personal computers), one or more special-purpose computers (e.g., devices specifically programmed to communicate with each other and/or the client system 105), or a combination of one or more general-purpose computers and one or more special-purpose computers. The host system 110 may be arranged to operate within or in concert with one or more other systems, such as, for example, one or more LANs ("Local Area Networks") and/or one or more WANs ("Wide Area Networks").

The host device 135 and host controller 140 each typically includes one or more hardware components and/or software components. An example of a host device 135 is a general-purpose computer (e.g., a personal computer) capable of responding to and executing instructions in a defined manner. Other examples include a special-purpose computer, a workstation, a server, a device, a component, other equipment or some combination thereof capable of responding to and executing instructions. An example of host controller 140 is a software application loaded on the host device 135 for commanding and directing communications enabled by the host device 135. Other examples include a program, a piece

of code, an instruction, a device, a computer, a computer system, or a combination thereof, for independently or collectively instructing the host device 135 to interact and operate as described herein. The host controller 140 may be embodied permanently or temporarily in any type of machine, component, equipment, storage medium, or propagated signal capable of providing instructions to the host device 135.

Fig. 2 illustrates a communication system 200 including a client system 205 communicating with a host system 210 through a communications link 215. Client system 205 typically includes one or more client devices 220 and one or more client controllers 225 for controlling the client devices 220. Host system 210 typically includes one or more host devices 235 and one or more host controllers 240 for controlling the host devices 235. The communications link 215 may include communication pathways 250, 255 enabling communications through the one or more delivery networks 260.

Examples of each element within the communication system of Fig. 2 are broadly described above with respect to Fig. 1. In particular, the host system 210 and communications link 215 typically have attributes comparable to those described with respect to host system 110 and communications link 115 of Fig. 1. Likewise, the client system 205 of Fig. 2 typically has attributes comparable to and illustrates one possible embodiment of the client system 105 of Fig. 1.

The client device 220 typically includes a general purpose computer 270 having an internal or external storage 272 for storing data and programs such as an operating system 274 (e.g., DOS, Windows™, Windows 95™, Windows 98™, Windows 2000™, Windows NT™, OS/2, or Linux) and one or more application programs. Examples of application programs include authoring applications 276 (e.g., word processing, database programs, spreadsheet programs, or graphics programs) capable of generating documents or other electronic content; client applications 278 (e.g., AOL client, CompuServe client, AIM client, AOL TV client, or ISP client) capable of communicating with other computer users, accessing various computer resources, and viewing, creating, or otherwise manipulating electronic content; and browser applications 280 (e.g., Netscape's Navigator or Microsoft's Internet Explorer) capable of rendering standard Internet content.

The general-purpose computer 270 also includes a central processing unit 282 (CPU) for executing instructions in response to commands from the client controller 225. In one

implementation, the client controller 225 includes one or more of the application programs installed on the internal or external storage 272 of the general-purpose computer 270. In another implementation, the client controller 225 includes application programs externally stored in and performed by one or more device(s) external to the general-purpose computer 270.

The general-purpose computer typically will include a communication device 284 for sending and receiving data. One example of the communication device 284 is a modem. Other examples include a transceiver, a set-top box, a communication card, a satellite dish, an antenna, or another network adapter capable of transmitting and receiving data over the communications link 215 through a wired or wireless data pathway 250. The general-purpose computer 270 also may include a TV ("television") tuner 286 for receiving television programming in the form of broadcast, satellite, and/or cable TV signals. As a result, the client device 220 can selectively and/or simultaneously display network content received by communications device 284 and television programming content received by the TV tuner 286.

The general-purpose computer 270 typically will include an input/output interface 288 for wired or wireless connection to various peripheral devices 290. Examples of peripheral devices 290 include, but are not limited to, a mouse 291, a mobile phone 292, a personal digital assistant 293 (PDA), a keyboard 294, a display monitor 295 with or without a touch screen input, a TV remote control 296 for receiving information from and rendering information to subscribers, and a video input device 298.

Although Fig. 2 illustrates devices such as a mobile telephone 292, a PDA 293, and a TV remote control 296 as being peripheral with respect to the general-purpose computer 270, in another implementation, such devices may themselves include the functionality of the general-purpose computer 270 and operate as the client device 220. For example, the mobile phone 292 or the PDA 293 may include computing and networking capabilities and function as a client device 220 by accessing the delivery network 260 and communicating with the host system 210. Furthermore, the client system 205 may include one, some or all of the components and devices described above.

Referring to Fig. 3, a communications system 300 is capable of delivering and exchanging information between a client system 305 and a host system 310 through a

communication link 315. Client system 305 typically includes one or more client devices 320 and one or more client controllers 325 for controlling the client devices 320. Host system 310 typically includes one or more host devices 335 and one or more host controllers 340 for controlling the host devices 335. The communications link 315 may include communication pathways 350, 355 enabling communications through the one or more delivery networks 360.

Examples of each element within the communication system of Fig. 3 are broadly described above with respect to Figs. 1 and 2. In particular, the client system 305 and the communications link 315 typically have attributes comparable to those described with respect to client systems 105 and 205 and communications links 115 and 215 of Figs. 1 and 2. Likewise, the host system 310 of Fig. 3 may have attributes comparable to and illustrates one possible embodiment of the host systems 110 and 210 shown in Figs. 1 and 2, respectively.

The host system 310 includes a host device 335 and a host controller 340. The host controller 340 is generally capable of transmitting instructions to any or all of the elements of the host device 335. For example, in one implementation, the host controller 340 includes one or more software applications loaded on the host device 335. However, in other implementations, as described above, the host controller 340 may include any of several other programs, machines, and devices operating independently or collectively to control the host device 335.

The host device 335 includes a login server 370 for enabling access by subscribers and routing communications between the client system 305 and other elements of the host device 335. The host device 335 also includes various host complexes such as the depicted OSP ("Online Service Provider") host complex 380 and IM ("Instant Messaging") host complex 390. To enable access to these host complexes by subscribers, the client system 305 includes communication software, for example, an OSP client application and an IM client application. The OSP and IM communication software applications are designed to facilitate the subscriber's interactions with the respective services and, in particular, may provide access to all the services available within the respective host complexes.

Typically, the OSP host complex 380 supports different services, such as email, discussion groups, chat, news services, and Internet access. The OSP host complex 380 is

generally designed with an architecture that enables the machines within the OSP host complex 380 to communicate with each other and employs certain protocols (i.e., standards, formats, conventions, rules, and structures) to transfer data. The OSP host complex 380 ordinarily employs one or more OSP protocols and custom dialing engines to enable access by selected client applications. The OSP host complex 380 may define one or more specific protocols for each service based on a common, underlying proprietary protocol.

The IM host complex 390 is generally independent of the OSP host complex 380, and supports instant messaging services irrespective of a subscriber's network or Internet access. Thus, the IM host complex 390 allows subscribers to send and receive instant messages, whether or not they have access to any particular ISP. The IM host complex 390 may support associated services, such as administrative matters, advertising, directory services, chat, and interest groups related to the instant messaging. The IM host complex 390 has an architecture that enables all of the machines within the IM host complex to communicate with each other. To transfer data, the IM host complex 390 employs one or more standard or exclusive IM protocols.

The host device 335 may include one or more gateways that connect and therefore link complexes, such as the OSP host complex gateway 385 and the IM host complex gateway 395. The OSP host complex gateway 385 and the IM host complex 395 gateway may directly or indirectly link the OSP host complex 380 with the IM host complex 390 through a wired or wireless pathway. Ordinarily, when used to facilitate a link between complexes, the OSP host complex gateway 385 and the IM host complex gateway 395 are privy to information regarding the protocol type anticipated by a destination complex, which enables any necessary protocol conversion to be performed incident to the transfer of data from one complex to another. For instance, the OSP host complex 380 and IM host complex 390 generally use different protocols such that transferring data between the complexes requires protocol conversion by or at the request of the OSP host complex gateway 385 and/or the IM host complex gateway 395.

Referring to Fig. 4, a communications system 400 is capable of delivering and exchanging information between a client system 405 and a host system 410 through a communication link 415. Client system 405 typically includes one or more client devices 420 and one or more client controllers 425 for controlling the client devices 420. Host

system 410 typically includes one or more host devices 435 and one or more host controllers 440 for controlling the host devices 435. The communications link 415 may include communication pathways 450, 455 enabling communications through the one or more delivery networks 460. As shown, the client system 405 may access the Internet 465 through the host system 410.

Examples of each element within the communication system of Fig. 4 are broadly described above with respect to Figs. 1-3. In particular, the client system 405 and the communications link 415 typically have attributes comparable to those described with respect to client systems 105, 205, and 305 and communications links 115, 215, and 315 of Figs. 1-3. Likewise, the host system 410 of Fig. 4 may have attributes comparable to and illustrates one possible embodiment of the host systems 110, 210, and 310 shown in Figs. 1-3, respectively. However, Fig. 4 describes an aspect of the host system 410, focusing primarily on one particular implementation of OSP host complex 480. For purposes of communicating with an OSP host complex 480, the delivery network 460 is generally a telephone network.

The client system 405 includes a client device 420 and a client controller 425. The client controller 425 is generally capable of establishing a connection to the host system 410, including the OSP host complex 480, the IM host complex 490 and/or the Internet 465. In one implementation, the client controller 425 includes an OSP application for communicating with servers in the OSP host complex 480 using exclusive OSP protocols. The client controller 425 also may include applications, such as an IM client application, and/or an Internet browser application, for communicating with the IM host complex 490 and the Internet 465.

The host system 410 includes a host device 435 and a host controller 440. The host controller 440 is generally capable of transmitting instructions to any or all of the elements of the host device 435. For example, in one implementation, the host controller 440 includes one or more software applications loaded on one or more elements of the host device 435. However, in other implementations, as described above, the host controller 440 may include any of several other programs, machines, and devices operating independently or collectively to control the host device 435.

The host system 410 includes a login server 470 capable of enabling communications with and authorizing access by client systems 405 to various elements of the host system 410, including an OSP host complex 480 and an IM host complex 490. The login server 470 may implement one or more authorization procedures to enable simultaneous access to the OSP host complex 480 and the IM host complex 490. The OSP host complex 480 and the IM host complex 490 are connected through one or more OSP host complex gateways 485 and one or more IM host complex gateways 495. Each OSP host complex gateway 485 and IM host complex gateway 495 may perform any protocol conversions necessary to enable communication between the OSP host complex 480, the IM host complex 490, and the Internet 465.

The OSP host complex 480 supports a set of services from one or more servers located internal to and external from the OSP host complex 480. Servers external to the OSP host complex 480 generally may be viewed as existing on the Internet 465. Servers internal to the OSP complex 480 may be arranged in one or more configurations. For example, servers may be arranged in centralized or localized clusters in order to distribute servers and subscribers within the OSP host complex 480.

In the implementation of Fig. 4, the OSP host complex 480 includes a routing processor 4802. In general, the routing processor 4802 will examine an address field of a data request, use a mapping table to determine the appropriate destination for the data request, and direct the data request to the appropriate destination. In a packet-based implementation, the client system 405 may generate information requests, convert the requests into data packets, sequence the data packets, perform error checking and other packet-switching techniques, and transmit the data packets to the routing processor 4802. Upon receiving data packets from the client system 405, the routing processor 4802 may directly or indirectly route the data packets to a specified destination within or outside of the OSP host complex 480. For example, in the event that a data request from the client system 405 can be satisfied locally, the routing processor 4802 may direct the data request to a local server 4804. In the event that the data request cannot be satisfied locally, the routing processor 4802 may direct the data request externally to the Internet 465 or the IM host complex 490 through the gateway 485.

The OSP host complex 480 also includes a proxy server 4806 for directing data requests and/or otherwise facilitating communication between the client system 405 and the Internet 465 through. The proxy server 4802 may include an IP ("Internet Protocol") tunnel for converting data from OSP protocol into standard Internet protocol and transmitting the data to the Internet 465. The IP tunnel also converts data received from the Internet in the standard Internet protocol back into the OSP protocol and sends the converted data to the routing processor 4802 for delivery back to the client system 405.

The proxy server 4806 also may allow the client system 405 to use standard Internet protocols and formatting to access the OSP host complex 480 and the Internet 465. For example, the subscriber can use an OSP TV client application having an embedded browser application installed on the client system 405 to generate a request in standard Internet protocol, such as HTTP ("HyperText Transport Protocol"). In a packet-based implementation, data packets may be encapsulated inside a standard Internet tunneling protocol, such as, for example, UDP ("User Datagram Protocol") and routed to the proxy server 4806. The proxy server 4806 may include a L2TP ("Layer Two Tunneling Protocol") tunnel capable of establishing a point-to-point protocol (PPP) session with the client system 405.

The proxy server 4806 also may act as a buffer between the client system 405 and the Internet 465, and may implement content filtering and time saving techniques. For example, the proxy server 4806 can check parental controls settings of the client system 405 and request and transmit content from the Internet 465 according to the parental control settings. In addition, the proxy server 4806 may include one or more caches for storing frequently accessed information. If requested data is determined to be stored in the caches, the proxy server 4806 may send the information to the client system 405 from the caches and avoid the need to access the Internet 465.

Referring to Fig. 5, a communications system 500 is capable of delivering and exchanging information between a client system 505 and a host system 510 through a communication link 515. Client system 505 typically includes one or more client devices 520 and one or more client controllers 525 for controlling the client devices 520. Host system 510 typically includes one or more host devices 535 and one or more host controllers 540 for controlling the host devices 535. The communications link 515 may include

communication pathways 550, 555 enabling communications through the one or more delivery networks 560. As shown, the client system 505 may access the Internet 565 through the host system 510.

Examples of each element within the communication system of Fig. 5 are broadly described above with respect to Figs. 1-4. In particular, the client system 505 and the communications link 515 typically have attributes comparable to those described with respect to client systems 105, 205, 305, and 405 and communications links 115, 215, 315, and 415 of Figs. 1-4. Likewise, the host system 510 of Fig. 5 may have attributes comparable to and illustrates one possible embodiment of the host systems 110, 210, 310, and 410 shown in Figs. 1-4, respectively. However, Fig. 5 describes an aspect of the host system 510, focusing primarily on one particular implementation of IM host complex 590. For purposes of communicating with the IM host complex 590, the delivery network 560 is generally a telephone network.

The client system 505 includes a client device 520 and a client controller 525. The client controller 525 is generally capable of establishing a connection to the host system 510, including the OSP host complex 580, the IM host complex 590 and/or the Internet 565. In one implementation, the client controller 525 includes an IM application for communicating with servers in the IM host complex 590 utilizing exclusive IM protocols. The client controller 525 also may include applications, such as an OSP client application, and/or an Internet browser application for communicating with the OSP host complex 580 and the Internet 565, respectively.

The host system 510 includes a host device 535 and a host controller 540. The host controller 540 is generally capable of transmitting instructions to any or all of the elements of the host device 535. For example, in one implementation, the host controller 540 includes one or more software applications loaded on one or more elements of the host device 535. However, in other implementations, as described above, the host controller 540 may include any of several other programs, machines, and devices operating independently or collectively to control the host device 535.

The host system 510 includes a login server 570 capable of enabling communications with and authorizing access by client systems 505 to various elements of the host system 510, including an OSP host complex 580 and an IM host complex 590. The login server 570 may

implement one or more authorization procedures to enable simultaneous access to the OSP host complex 580 and the IM host complex 590. The OSP host complex 580 and the IM host complex 590 are connected through one or more OSP host complex gateways 585 and one or more IM host complex gateways 595. Each OSP host complex gateway 585 and IM host complex gateway 595 may perform any protocol conversions necessary to enable communication between the OSP host complex 580, the IM host complex 590, and the Internet 565.

To access the IM host complex 590 to begin an instant messaging session, the client system 505 establishes a connection to the login server 570. The login server 570 typically determines whether the particular subscriber is authorized to access the IM host complex 590 by verifying a subscriber identification and password. If the subscriber is authorized to access the IM host complex 590, the login server 570 employs a hashing technique on the subscriber's screen name to identify a particular IM server 5902 for use during the subscriber's session. The login server 570 provides the client system 505 with the IP address of the particular IM server 5902, gives the client system 505 an encrypted key (i.e., a cookie), and breaks the connection. The client system 505 then uses the IP address to establish a connection to the particular IM server 5902 through the communications link 515, and obtains access to that IM server 5902 using the encrypted key. Typically, the client system 505 will be equipped with a Winsock API ("Application Programming Interface") that enables the client system 505 to establish an open TCP connection to the IM server 5902.

Once a connection to the IM server 5902 has been established, the client system 505 may directly or indirectly transmit data to and access content from the IM server 5902 and one or more associated domain servers 5904. The IM server 5902 supports the fundamental instant messaging services and the domain servers 5904 may support associated services, such as, for example, administrative matters, directory services, chat and interest groups. In general, the purpose of the domain servers 5904 is to lighten the load placed on the IM server 5902 by assuming responsibility for some of the services within the IM host complex 590. By accessing the IM server 5902 and/or the domain server 5904, a subscriber can use the IM client application to view whether particular subscribers ("buddies") are online, exchange instant messages with particular subscribers, participate in group chat rooms, trade files such

as pictures, invitations or documents, find other subscribers with similar interests, get customized news and stock quotes, and search the Web.

In the implementation of Fig. 5, the IM server 5902 is directly or indirectly connected to a routing gateway 5906. The routing gateway 5906 facilitates the connection between the IM server 5902 and one or more alert multiplexors 5908, for example, by serving as a link minimization tool or hub to connect several IM servers to several alert multiplexors. In general, an alert multiplexor 5908 maintains a record of alerts and subscribers registered to receive the alerts.

Once the client system 505 is connected to the alert multiplexor 5908, a subscriber can register for and/or receive one or more types of alerts. The connection pathway between the client system 505 and the alert multiplexor 5908 is determined by employing another hashing technique at the IM server 5902 to identify the particular alert multiplexor 5908 to be used for the subscriber's session. Once the particular multiplexor 5908 has been identified, the IM server 5902 provides the client system 505 with the IP address of the particular alert multiplexor 5908 and gives the client system 505 an encrypted key (i.e., a cookie). The client system 505 then uses the IP address to connect to the particular alert multiplexor 5908 through the communication link 515 and obtains access to the alert multiplexor 5908 using the encrypted key.

The alert multiplexor 5908 is connected to an alert gate 5910 that, like the IM host complex gateway 595, is capable of performing the necessary protocol conversions to form a bridge to the OSP host complex 580. The alert gate 5910 is the interface between the IM host complex 590 and the physical servers, such as servers in the OSP host complex 580, where state changes are occurring. In general, the information regarding state changes will be gathered and used by the IM host complex 590. However, the alert multiplexor 5908 also may communicate with the OSP host complex 580 through the IM gateway 595, for example, to provide the servers and subscribers of the OSP host complex 580 with certain information gathered from the alert gate 5910.

The alert gate 5910 can detect an alert feed corresponding to a particular type of alert. The alert gate 5910 may include a piece of code (alert receive code) capable of interacting with another piece of code (alert broadcast code) on the physical server where a state change occurs. In general, the alert receive code installed on the alert gate 5910 instructs the alert

broadcast code installed on the physical server to send an alert feed to the alert gate 5910 upon the occurrence of a particular state change. Upon detecting an alert feed, the alert gate 5910 contacts the alert multiplexor 5908, which in turn, informs the client system 505 of the detected alert feed.

5 In the implementation of Fig. 5, the IM host complex 590 also includes a subscriber profile server 5912 connected to a database 5914 for storing large amounts of subscriber profile data. The subscriber profile server 5912 may be used to enter, retrieve, edit, manipulate, or otherwise process subscriber profile data. In one implementation, a subscriber's profile data includes, for example, the subscriber's buddy list, alert preferences,
10 designated stocks, identified interests, and geographic location. The subscriber may enter, edit and/or delete profile data using an installed IM client application on the client system 505 to interact with the subscriber profile server 5912.

Because the subscriber's data is stored in the IM host complex 590, the subscriber does not have to reenter or update such information in the event that the subscriber accesses
15 the IM host complex 590 using new or a different client system 505. Accordingly, when a subscriber accesses the IM host complex 590, the IM server 5902 can instruct the subscriber profile server 5912 to retrieve the subscriber's profile data from the database 5914 and to provide, for example, the subscriber's buddy list to the IM server 5902 and the subscriber's alert preferences to the alert multiplexor 5908. The subscriber profile server 5912 also may
20 communicate with other servers in the OSP host complex 590 to share subscriber profile data with other services. Alternatively, user profile data may be saved locally on the client device 505.

Referring to Fig. 6, a sender 602a, a recipient 602b, and a host 604 interact according to a procedure 600 to transfer audio data. The procedure 600 may be implemented by any
25 suitable type of hardware, software, device, computer, computer system, equipment, component, program, application, code, storage medium, or propagated signal.

Examples of each element of Fig. 6 are broadly described above with respect to Figs. 1-5. In particular, the sender 602a and the recipient 602b typically have attributes comparable to those described with respect to client devices 120, 220, 320, 420, and 520
30 and/or client controllers 125, 225, 325, 425, and 525. The host 604 typically has attributes comparable to those described with respect to host device 135, 235, 335, 435, and 535 and/or

host controllers 140, 240, 340, 440, and 540. The sender 602a, the recipient 602b, and/or the host 604 may be directly or indirectly interconnected through a known or described delivery network.

5 The sender 602a and the recipient 602b are each associated with a subscriber. To allow file transfers, each subscriber sets certain preferences for permitting files to be transferred to and from other subscribers. For example, the sender and recipient may identify screen names of subscribers who have permission to send files to them or retrieve files from them. Typically, each subscriber will be presented with a graphical user interface that permits selection among various transfer preferences. A subscriber's transfer preferences
10 may be maintained locally at the client or remotely at the host 604.

In general, the sender 602a and the recipient 602b communicate over an open connection, such as an open TCP connection established through the host 604. Typically, the sender 602a and the recipient 602b each include a Winsock API for establishing an open TCP connection to the host 604 and a client application for accessing the host 604. The
15 sender 602a and the recipient 602b connect to the host 604 to establish the connection.

The sender 602a and the recipient 602b use the connection to communicate with the host 604 and with each other. The connection remains open during the time that the sender 602a and the recipient 602b are accessing the host 604. To access the host 604, the sender 602a and the recipient 602b each send a separate request to the host 604. The request
20 identifies the associated subscriber to the host 604 and to other subscribers using a unique screen name. The host 604 verifies a subscriber's information (e.g., screen name and password) against data stored in a subscriber database. If the subscriber's information is verified, the host 604 authorizes access. If the subscriber's information is not verified, the host 604 denies access and sends an error message.

25 Upon accessing the host 604, a "buddy list" is displayed to the subscriber. In general, a subscriber's buddy list is a user interface that lists the online status and capabilities of certain screen names, i.e., "buddies", identified the subscriber. In particular, the host 604 informs the sender whether identified buddies are online, i.e., currently accessing the host 604. The host 604 also informs any subscriber who has identified the sender as a buddy that
30 the sender is currently online. The buddy list also facilitates instant messaging communication between subscribers. A subscriber can activate an instant messaging

message user interface pre-addressed to a buddy simply by clicking the screen name of a buddy on the buddy list. If a recipient is not a "buddy," the first subscriber must activate a blank instant messaging user interface and then address the instant message to the screen name of the intended recipient. When necessary, a subscriber can look up the screen name of an intended recipient using the intended recipient's e-mail address.

In addition to exchanging instant messages with online buddies, the sender may participate in group chat rooms, locate other subscribers with similar interests, get customized news and stock quotes, search the Web, and transfer files to and from other subscribers. In one implementation, a sender 602a, a recipient 602b, and a host 604 interact according to a procedure 600 to transfer audio data.

The transfer of audio data extends the functionality of instant messaging by allowing the sender 602a and the recipient 602b to communicate peer to peer via audio, i.e., microphone and speaker. In one implementation, the sender initiates the process 600 by designating one or more recipients to receive an instant message (e.g., a text message). If the intended recipients are "buddies" of the sender 602a, the sender 602a may confirm the online status and capabilities of each recipient prior to sending the video message by viewing the "buddy list." After a subscriber composes an instant message and clicks a SEND button, the instant message is sent from the sender 602a to the host (step 605).

After receiving the instant message from the sender 602a, the host 704 authenticates the instant message (step 610). In addition to the textual body, the instant message may include header information identifying the message type, the screen name and/or IP address of the sender and recipient, and a randomly generated security number. The instant message may be authenticated by, for example, using a reverse look-up table to match the screen names and/or IP addresses with those of valid subscribers. In the event that either the sender 602a or the recipient 602b is not associated with a valid subscriber, the host 604 reports an error message.

Once the instant message is verified, the host 604 determines the capabilities of the recipient (step 615). For example, the host 604 may monitor and update the online status, client version, and device type of all connected subscribers in real time. The capability to receive audio data may depend on hardware (e.g., device type), software (e.g., client version), and/or transfer preferences (e.g., blocked screen names). To be talk enabled, both

the talk software and audio equipment must be available. The host 604 then reports the capabilities of the recipient to the sender (step 620).

Upon receiving the report from the host 604, the sender 602a displays a UI according to the capabilities of the sender and/or the recipient 702b (step 625). If the sender 602a is not talk enabled, then a standard instant messaging user interface is displayed. If the sender 602a is talk enabled, but the recipient 602b is not talk enabled, a START TALK UI having a grayed START TALK button is displayed. If both the sender 602a and the recipient 602b are talk enabled, a START TALK UI having a functioning START TALK button is displayed.

The process 600 continues with the host 604 sending the instant message to the recipient 602b (step 630). The recipient 602b accepts the initial text message from the host 604 (step 635) and displays a UI according to the capabilities of the sender 602a and/or the recipient 602b (step 640). If the recipient 602b is not talk enabled, then a standard instant messaging UI is displayed. If the recipient 602b is talk enabled, but the sender 602a is not talk enabled, an instant messaging UI having a grayed START TALK button is displayed. If both the recipient 602b and the sender 602a are talk enabled, an instant messaging UI with a functioning START TALK button is displayed.

If both sides are talk enabled, both the sender 602a and the recipient 602b have a START TALK UI displayed. When the START TALK UI is displayed, a subscriber can initiate a talk session. In one implementation, the sender 602a initiates a talk session by sending a talk request to the host 604 (step 645). The talk request may contain information including, but not limited to, the message type, the screen name and/or IP address of the sender and recipient, and a randomly generated security number. When the sender 602a clicks the START TALK UI, the START TALK UI transitions to an END TALK UI.

Upon receiving the talk request, the host 604 authenticates the talk request from the sender 602a (step 650). The host 604 may authenticate the talk request by, for example, using a reverse look-up table to match the screen names and/or IP addresses with those of valid subscribers. In the event that either the sender 602a or the recipient 602b is not associated with a valid subscriber, the host 604 reports an error message.

After verifying the talk request, the host 604 sends the talk request to the recipient 602b (step 655). Upon receiving the talk request, the START TALK UI displayed by the recipient 602b transitions to a CONNECT UI (step 660). The CONNECT UI informs the

recipient 602b that the sender 602a wants to engage in a talk session. At this point, the recipient 602b may ignore the talk request, accept the talk request, or terminate the instant message session.

If the recipient 602b accepts the talk request by clicking the CONNECT UI (step 665), the CONNECT UI transitions to the END TALK UI and the host 604 establishes a talk session (step 670). When a talk session is active, users can talk to each other. At this point, END TALK UI is displayed by both the sender 602a and the recipient 602b. The talk session (steps 675a-b) remains active until one of the users clicks END TALK UI. After one of the users clicks the END TALK UI, both the sender 602a and the recipient 602b will display the START TALK UI, allowing either side to initiate yet another talk session.

If the sender 602a disengages from the talk session before the recipient connects, the CONNECT UI at the recipient 602b transitions back to the START TALK UI. If both users click the START TALK UI simultaneously, the host will ignore one of the START TALK clicks such that one user will display the END TALK UI and the other will display the CONNECT UI. If the sender clicks the START TALK UI prior to the recipient 602b accepting the initial text message, the recipient 602b does not display the START TALK UI, but instead immediately displays the CONNECT UI.

In one implementation, a talk tool establishes an active talk session using three communication channels: a Generic Signaling Interface (GSI) channel, a control channel, and an audio channel. The talk tool uses the GSI channel to establish the initial connection. During this connection, the local IP addresses are exchanged. After the initial connection phase is done, the GSI channel is no longer used. By using the GSI channel, the exchange of local IP addresses is only done when both users permit such an exchange, i.e., by clicking on the CONNECT UI. These actions protect users from having their local IP addresses automatically obtained without their consent.

The control channel is a TCP/IP socket, for which the IP address and port number of the remote side are obtained through the GSI channel. The control channel is used to send/receive control attributes of the talk session while the session is active. For example, because some firewalls will not allow an external connection to a socket on the inside of the firewall, the talk tool attempts a connection from both sides of the session. This action allows a connection to be made if there is a maximum of one firewall within the connection. If there

is a firewall on both sides, the chances are that no connection can be made and the talk session will fail. To work across two firewalls, the user must obtain the port range used by talk such that one of the firewalls can be modified to permit the range to pass through the firewall.

5 The audio channel is a TCP/IP socket used to transport audio packets. This channel can either be UDP or TCP. In general, UDP is used since it minimizes latency. However, because some firewalls will not pass through UDP packets, the audio channel may have to use TCP. The talk tool indicates the mode (i.e., TCP, UDP), or employs an auto mode in which the talk tool attempts a UDP test and resorts to TCP upon failure of UDP.

10 Talk sessions may work in either full or half duplex. Full duplex is when both users can talk at the same time. Half duplex is where only one user can talk at a time. A client device is determined to be incapable of handling full duplex, for example, if the CPU is too slow to compress/decompress audio simultaneously and/or the microphone and speakers cannot be opened simultaneously. If a client device is marked as half duplex, then any talk
15 session used by that client device becomes a half duplex session, regardless of whether another device can handle duplex mode. In one implementation, a TALK/LISTEN button on the END TALK UI supports half duplex operation. This button has two states: LISTEN or TALK. If the talk session is full duplex, this button is not shown. If the button reads TALK at both the sender 702a and the recipient 702b (Initial Half Duplex), the first user to click
20 TALK is allowed to talk and the other user is forced to listen. The user who is listening has a grayed out TALK button (Half Duplex Listen) and the user who is talking has a LISTEN button (Talking Half Duplex). When the LISTEN button is clicked, the user who is talking allows the user who is listening to talk.

25 The talk tool that enables the audio transfer (talk) functionality may be any type of client controller (e.g., software, application, program) loaded on to a client device. The talk tool supports use by different OSP and IM clients. The talk tool is responsible for responding to user interfaces and translating user commands into the appropriate actions with the client device. For example, the talk tool opens, reads, writes, and closes the physical components on the client devices needed for audio. The talk tool also controls audio and
30 control channels with callbacks being executed to indicate status change. When the talk tool is loaded, the talk tool determines if the client device is capable of handling full duplex.

The talk tool also may allow the user to control the volume for the speaker and microphone. In one implementation, the user speaks into a microphone and the audio data are recorded into memory. While in the record mode, the average level of the speaker's voice is indicated on a level meter displayed on a user interface of the talk tool. A slider control is used to adjust the input level to an optimal value. After the speaker stops speaking, the speaker's stored speech is played back through the computer's audio output device. The speaker level slider control may be used to adjust the output level to an acceptable volume. If the user starts to speak again, the talk tool reverts to the record mode and the cycle repeats. Once the user is satisfied with the settings, the user can save the settings for use in subsequent talk sessions.

The talk tool may support additional functionality including, but not limited to, multi-conferencing, hold, and muting. Multi-conferencing allows more than two users to engage in a talk session. Hold allows the suspension of an active talk session in order to connect to another talk session. Muting turns off the microphone to prevent user feedback/echo during full duplex mode.

The talk tool also may include security features to protect the integrity of transferred data. For example, the talk tool may compress data using a proprietary algorithm or may send the data in a proprietary protocol. To further improve security, the talk tool may select the port numbers at random from a large range.

In general, an instant messaging talk session is similar to a telephonic session in that it has the same three states: not connected (hung up), connecting (ringing), and connected (talking). As described above, these states and the ability to switch among them are supported by corresponding UIs, namely a START TALK UI (not connected), a CONNECT UI (ringing), and an END TALK UI (connected).

Fig. 7 illustrates one example of a START TALK UI. As shown in Fig. 7, a START UI 700 includes an instant message box 705 having a START TALK button 710 for requesting a talk session.

Fig. 8 illustrates one example of a CONNECT UI. As shown in Fig. 8, a UI 800 includes an instant message box 805 having a CONNECT button 810 for accepting a request to initiate a talk session.

Fig. 9 illustrates one example of an END TALK UI. As shown in Fig. 9, a UI 900 includes an instant message box 905 having an END TALK button 910 for terminating a talk session.

Fig. 10 illustrate one example of a half duplex user interface. As shown in Fig. 10, a
5 UI 1000 includes an instant message box 1005 having a TALK button 1010. The bottom
1010 is greyed out or otherwise disabled when the other party is talking.

Other embodiments are within the scope of the following claims.